

AD-A247 049 DOCUMENTATION PAGE

Form Approved  
OMB No. 0704-0188

Information is furnished to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and reviewing the collection of information, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this form including this burden estimate, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Avenue, 4th Floor, Washington, DC 20540, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 1/10/92		3. REPORT TYPE AND DATES COVERED Annual Technical- 12/1/90 to 12/31/91	
4. TITLE AND SUBTITLE Neuropsychological Components of Object Identification				5. FUNDING NUMBERS NR-61102F G:AFOSR-91-01100	
6. AUTHOR(S) Stephen M. Kosslyn, Ph.D.				PR: 2313 TA: BS	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Harvard University-Department of Psychology 33 Kirkland Street Cambridge, MA 02138				8. PERFORMING ORGANIZATION REPORT NUMBER AEOSR-TR- 92 01361	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR Bolling Air Force Base Building 410 Washington, DC 20332				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES DTIC ELECTE MAR 05 1992 S D D					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release:Distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Four kinds of investigations of the structure of high-level visual processing were conducted during the past year. First, we carried out case studies of individual brain-damaged patients, finding evidence that curved edges are processed separately from straight edges, that location information sometimes can be used to encode some characteristics of shape, and that a decrease in overall "activation" level can selectively impair performance on some tasks. Second, we administered a set of 27 tasks to a group of 17 brain-damaged patients; these tasks were designed to assess the efficacy of specific subsystems. We have preliminary evidence that most (89%) of the subsystems double dissociate—suggesting that these subsystems are in fact distinct. Third, we constructed computer simulation models to explore properties of the high-level visual system, and found support for the distinction between subsystems that compute two distinct kinds of spatial relations. Finally, some of the tasks we had developed to study deficits in brain-damaged patients were used to study the visual-spatial abilities of air force pilots; we found that pilots are particularly good at mental rotation and encoding metric distance information.					
14. SUBJECT TERMS Neuropsychology, vision, computational models.				15. NUMBER OF PAGES 33	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT (u)		

*Neuropsychological Components of Object Identification*  
1991 Annual Report  
S. M. Kosslyn, PI

We have conducted four kinds of investigations during the past year to study the structure of high-level visual processing. First, we have conducted case studies of individual brain-damaged patients, finding evidence for specific properties of the intact processing system. Second, we have administered a set of 27 tasks to a group of 17 brain-damaged patients, and are analyzing the data for double dissociations, which would provide evidence for the existence of distinct processing subsystems. Third, we have constructed computer simulation models to explore further properties of the visual system. Finally, it occurred to us that some of the tasks we had developed to study deficits in brain-damaged patients might also have a use in explaining superior visual-spatial abilities. Thus, we tested a group of air force pilots on some of these tasks. Each type of work is summarized below.

*Case Studies of Brain-Damaged Patients*

In the 19th century researchers observed behavioral deficits following brain damage, and inferred that the processes that normally would have produced the disrupted behavior were damaged. These sorts of inferences do not necessarily follow, however (e.g., see Gregory, 1969; Kosslyn & Van Kleeck, 1990). Logically, behavioral dysfunction can reflect not only damaged processing subsystems, but also disconnections between subsystems (cf. Geschwind, 1966), lowered "activation" levels, and various types of compensations (for a review and discussion of such ideas, see Kosslyn & Koenig, 1992). Nevertheless, by far the majority of studies of brain-damaged patients today have a decidedly 19th-century flavor: Researchers assume that the damage has affected discrete representations or processes, and the resulting behavior reflects the operation of an otherwise intact system—which does not compensate or act any differently than it did prior to damage (Caramazza, 1984, calls this the "transparency assumption"). After Kosslyn & Van Kleeck (1990) pointed this out, Caramazza (in press) rose to defend the traditional approach—and Kosslyn & Intriligator (in press) then considered carefully Caramazza's responses and found them less than compelling.

This debate, however, has remained at a fairly abstract level. We have conducted a series of case studies to demonstrate the power of a more sophisticated approach toward understanding the causes of behavioral dysfunction following brain damage. I will briefly summarize here three case studies, which focus on different types of underlying etiologies.

*Patient G.A.: A selective deficit for curvature*

The first patient was studied within the traditional framework. We used a series of converging measures to make the case that the deficit really did reflect the selective dysfunction of a specific type of process or a disconnection of its input to "higher" processes.

Many different kinds of information—not only shape, color, and texture, but also distinctive patterns of movement (e.g., see Farah, 1990; Kosslyn & Koenig, 1992)—are used to recognize and identify objects. In this study we showed that shape itself is not encoded by a unitary system. Rather, it appears that information about



curvature is at some point processed separately from information about straight lines and contours.

The idea that the perception of curvature can be selectively disrupted is interesting in part because distinct "curvature detector" cells have never been reported in visual cortex. A possible hint of the solution of this mystery was offered by Lehky and Sejnowski (1988), who described a neural network model that was trained to compute the shapes of curved objects from shading information. After being trained to perform this task, Lehky and Sejnowski discovered that the network had developed "hidden units" that responded when the termini of line segments were presented as stimuli. Such "end-stopped" cells have been reported in monkey visual cortex, and it is possible that sets of such cells work together to compute curvature. Our finding that the perception of curvature can be selectively disrupted suggests that the end-stopped cells or other neurons responsible for this computation may have a special property that makes them particularly vulnerable to damage caused by anoxia, such as occurred in our patient.

**Subjects.** At the time of testing patient G.A. was 13 years, 11 months old. He is a white male who sustained a cardiopulmonary arrest with subsequent coma in an apparently "routine" fall from a crib at 1 year, 5 months. He is right handed, and has no family history of non-right handedness, learning disorders, or medical, neurological, or emotional disorders. Evaluation of visual function in 1988 by Dr. Leonard Scinto of the Applied Science Laboratories documented difficulty in directing gaze to specified points in the visual field, impaired visual pursuit, unusual saccades, a puzzling selective allocation of gaze to either left or right visual space, visual search scan paths characterized by tight fixation clusters (which were markedly mitigated when he verbalized his scanning strategy), and apparently aimless wandering of gaze around central fixation marks in the course of specific visual tasks. These behaviors are consistent with some of the symptoms of Balint's syndrome, and the overall pattern of performance on both neurological and neuropsychological examination has been interpreted as consistent with bilateral "watershed area" lesions. However, an MRI taken at the time of testing was read as normal, and previous CT scans showed no evidence of localized damage.

G.A. scored within the average range for adults on the "Facial Recognition task" of Benton, Hamsher, Varney, and Spreen (1978, 1983), but only following strikingly extended response latencies as he compared the possible response choices to the target stimulus on a feature-by-feature basis. He also has trouble tracking multiple related components in a story, lack of appreciation of elapsed time, makes perceptual errors on the Boston Naming Test (Kaplan, Goodglass and Weintraub, 1983), and mislocates elements in the delayed Rey-Osterrieth Complex Figure test. However, he performs well on the Wisconsin Card Sort task (Heaton, 1981) and has average performance on the Stroop Color Word test (Golden, 1978). On motor tasks he prefers the right side, but is consistently more competent with the left.

At the time of our testing, G.A.'s visual behavior was notable for the following: First, he could not encode facial features as uniquely representative of a given person and could not recognize familiar persons by sight; he can recognize a face as a face, but cannot recognize family members, his own photograph, or regularly seen acquaintances. He asked one clinician to "say something so I can tell if I know you." He has been markedly confused over the years by his mother's changes of hairstyle or coloring. Second, he cannot visually identify subtly distinct members of other sets, such as his tent at camp one year. Third, he does not scan faces or other



Dist	even and/or Special
A-1	

complex objects systematically. Nevertheless, he can identify objects seen across the room and reads road signs and the like fluently. Fourth, he appears to recognize persons if they are moving, suggesting that he can encode distinctive patterns of movement.

In addition, we tested 11 of G.A.'s classmates at their private school, who volunteered to participate as paid subjects (and had parental permission to do so). These subjects had no documented medical or neurological problems, and ranged from 11 to 14 years of age; three were male and eight were female. There was one left-handed male subject, and all remaining subjects were right handed.

*Experiment 1.* We began by exploring the most basic ability to distinguish curvature, namely perceptual discrimination. On each trial of this experiment, the subjects saw a standard form, which was either a curved contour or an angular approximation to such a contour, and were to select that shape from a set of five alternatives. The curved and angular shapes could be relatively simple or complex, depending on the number of inflection points. Separate analyses were performed for response times and error rates. Each type of data was analyzed in two ways. First, we performed analyses of variance to compare G.A.'s performance with that of the control group. These analyses used stimulus items as the random effect, which allows us to generalize over items. (We computed a mean for each item for the control subjects, and analyzed these means with G.A.'s scores for each item.) Second, in order to discover whether the results would generalize over subjects, we computed the .05 (two-tailed) confidence intervals around the means for the subjects and for G.A., and noted whether these confidence intervals overlapped; if the two corresponding confidence intervals did not overlap, we concluded that G.A. had a deficit. All deficits noted below were significant in both measures.

G.A. did require more time for curved stimuli (5296 and 4538 ms for curved and angular stimuli for G.A., compared to 3083 and 3085 ms for the control subjects). Indeed, G.A.'s deficit could be isolated to the complex curved stimuli, which required 1489 ms more than the complex angular stimuli, whereas the simple curved stimuli required only 28 ms more than the simple angular ones. We analyzed the error rates in the same way, and found no hints of any effects or interactions involving stimulus type.

*Experiment 2.* We were concerned that the difference in response times might simply reflect G.A.'s comparison strategy; perhaps he spends more time looking back and forth between the standard and probe stimuli, and compares relatively few angles for the straight stimuli versus relatively many points for the curved stimuli. In the previous experiment, G.A. needed to compare a "target" stimulus to a set of probe stimuli. This task requires one to remember the standard (or parts of it) long enough to make the comparison. It is possible that G.A.'s problem is only in encoding the stimulus, but once it is encoded it is stored in memory normally. Sternberg (1966) argues that an encoding stage is independent from subsequent storage and comparison stages, which operate in the same way regardless of the amount of time needed to encode the stimulus. But it is also possible that G.A. not only encodes curvature more slowly, but also less effectively—and so this information is not as usefully represented in memory.

Experiment 1 was modified in only one way: Each trial was divided into two parts, with the first part containing only the target and the second containing only the five alternatives, at which point the subjects responded. G.A. again required more time for curved stimuli (curved and angular stimuli required 3349 and 2796

ms for G.A., compared to 2088 and 2148 ms for the control subjects). Unlike in Experiment 1, G.A. required more time for curved stimuli than angular stimuli for both the complex and simple stimuli. Again, however, there was no difference in accuracy between G.A. and the control subjects for the different stimulus types.

We next considered the results from Experiments 1 and 2 together to discover whether the effect of curvature was more extreme in Experiment 2. Although we again found evidence that G.A. was disproportionately slower for the curved stimuli than for the angular ones, relative to the controls, we found no evidence that this interaction was different in the two experiments. Thus, it seems likely that G.A.'s deficit involves not memory for curvature but a difficulty in perceptually encoding curvature.

*Experiment 3.* Both Experiments 1 and 2 required the subjects to encode and compare shapes. If G.A. has a deficit in perceptual encoding per se, it should be evident even when shapes need not be compared. In this experiment we consider a task in which G.A. does not compare shapes, but merely encodes them in the course of making a simple perceptual judgment. In this experiment we investigated the effects of curvature in a very simple shape-encoding task. Subjects were asked merely to decide whether a black X mark was on or off a gray geometric pattern. These patterns were composed of straight edges or curved edges. The straight-edged patterns were constructed by filling in contiguous cells of a 4x5 grid, and the curved patterns were formed by then smoothing the corners and edges of these patterns. The X marks were made of straight lines for straight-edged patterns, and curved lines for curved patterns.

The results revealed that G.A. required more time to evaluate curved stimuli than straight stimuli (967 versus 591 ms), whereas the control subjects required a similar amount of time in both cases. Again, there was no difference in error rates.

*Experiment 4.* We reasoned that if G.A. has difficulty encoding curved lines, then placing random curved line fragments over the curved patterns used in Experiment 3 should selectively impair his performance, compared to the effects of adding corresponding straight lines to the straight patterns. And in fact, we again found that G.A. required more time for curved versus straight stimuli, relative to the control subjects (for G.A., the means were 1113 and 678 ms for curved and straight stimuli, whereas for the control subjects they were 953 and 890 ms). We now also found that G.A. made more errors for curved stimuli relative to the control subjects (he made 16.7 and 4.2% errors for curved and straight stimuli, compared to 2.3 and 2.1% errors for the control subjects).

It is of interest that G.A. showed increased error rates in this experiment, when curved line segments were imposed on a stimulus. All of the previous experiments revealed that he required more time to compare curved lines, but not that he was less accurate. It is possible that the combination of a curved shape and curved "noise" simply overloaded G.A.'s capacity to encode curvature, leading to errors.

*Experiment 5.* If G.A. has a deficit in encoding curvature, this should be evident even when he views more naturalistic stimuli. To investigate this possibility, we showed our subjects a series of pictures of common objects, each of which was followed by a word. The subjects were to decide, as quickly as possible, whether the word was a correct name for the object. Half of the objects had predominantly straight edges, and half had predominantly curved edges. In addition, we placed random lines over some of the pictures, and examined the effect

of straight versus curved noise on verification times. Moreover, half of each type of picture were degraded; we thought that if G.A. has difficulty encoding curvature, this should be especially apparent when the contours are degraded—as often happens in natural environments, where smoke, occluding objects, shadows, and so on obscure parts of images.

We found that G.A. has selective difficulty in processing curved pictures, particularly when they are degraded. In addition, the confidence intervals suggested that G.A. made more errors than the control subjects when curved noise was present.

*Experiment 6.* Finally, we reasoned that if G.A. does in fact have a problem in encoding curvature, it should affect his ability to encode curved lines in general—including when he reads. Thus, we examined his ability to read text in a highly curved script font compared to a less curved, standard Roman typeface. And in fact, in a fixed time period, G.A. read fewer words in the script font relative to the Roman font than did the control subjects (with means of 24.6 vs. 36.1 for script and roman, compared to 28.6 and 36.1 for the control subjects). Although the effect was relatively small, it was very consistent and statistically significant. This finding is striking because G.A. is a relatively fluent reader, and has had much experience with typefaces like the ones we used. It is also of some interest that, in general, all subjects generally read the script font more slowly and less accurately. If this difference persists into adulthood, script or ornate fonts should be used with care when readers are expected to read particularly quickly or accurately.

In summary, the results show strong convergence: We have found a brain-damaged patient who performs poorly in a variety of tasks involving stimuli composed of curved lines. This is true when he must compare a curved line with other curved lines, whether the standard is present or must be remembered, when he examines curved shapes to determine whether an X is on or off, when he names degraded curved objects, and when he reads text in a curved script font. Moreover, his performance is slowed down when curved line segments are placed over a curved pattern in some circumstances but not others; we found such a decrement when and he was searching for a curved X target on a curved pattern, but only partial evidence (in the confidence intervals) when he was verifying the names of pictures.

The simplest account for these findings is that G.A. has difficulty encoding curvature. This inference is intriguing in light of his most striking clinical deficit, prosopagnosia. Although we cannot establish a causal relation between the two deficits, it may not be an accident that a person who has trouble encoding curvature also has difficulty encoding faces.

Most of the deficits we assessed were reflected in increased response times rather than more errors. This finding is consistent with the results of his prior neuropsychological testing; G.A. could recognize pictures of faces, but he required abnormal amounts of time. These results are in sharp contrast to G.A.'s performance in natural settings, where he has trouble recognizing faces even when he is given as much time as he wants to examine them. It is of interest, then, that G.A. did make more errors than the control subjects for curved stimuli in Experiment 4, which arguably incorporated the most complex stimuli. Three-dimensional faces are composed of a large number of curved contours, and these stimuli might overload his ability to encode curvature. We suspect, however, that if a deficit in encoding curvature is at the root of G.A.'s prosopagnosia, this will not be

true for all such patients; there are probably many ways of disrupting processing that will affect the encoding of faces (e.g., see Chapter 3 of Kosslyn & Koenig, 1992).

*Patient R.V.: "Automatic" compensations*

Our conception of information processing in the brain rests heavily on the idea of concurrent processing: Not only is information processed in cascade (with "pipelining" taking place), but also numerous "channels" and "streams" are operating in parallel (Kosslyn & Koenig, 1992). If so, then damaging one process will not cause merely local disruption, but will affect dynamic interactions in the system more generally. We examined one consequence of this idea by studying R.V., a patient with focal damage to the left frontal lobe. This damage disrupted the uncinate fasciculus and the inferior longitudinal fasciculus, thereby partially de-energizing parts of posterior cortex. Indeed, positron emission tomography revealed reduced metabolism in the occipital-temporal area, which is known to be involved in visual encoding.

We conducted many experiments on R.V., which were summarized in the previous grant's Final Report. Upon considering them further, we realized that their results collectively pointed towards an interesting new idea. Thus, we reanalyzed some of the data, exploring the possibility that brain damage could have the effect of altering the relative speeds of different processes. The results appear to show that behavior can be disrupted when the components of a system no longer interact normally, even though the components generally can still perform their individual functions. Our investigations focus on "high-level" visual processes (such as those used to identify objects), which make use of stored information, in contrast to low-level processes (such as those used to detect sharp changes in intensity), which are driven purely by the stimulus input. High-level visual processing depends in large part on two major pathways (e.g., Levine, 1982; Ungerleider & Mishkin, 1982). The ventral pathway runs from the occipital lobe to areas in the inferior temporal lobe; this system encodes "object properties," such as shape. The dorsal pathway, in contrast, runs from the occipital lobe up to areas in the parietal lobe; this system encodes "spatial properties," such as location.

Our experiments allowed us to explore the relation between the processing of object properties and spatial properties. Logically, a shape is nothing more than a juxtaposition of points at specific locations. Thus, under some circumstances it should be possible to use spatial processing to encode shapes, but in an unusual way—as a collection of occupied locations. We reasoned that if the ventral, object-properties encoding system is impaired, then the dorsal system may in some circumstances compensate by encoding shapes; in this case, however, the final behavior should depend on the ease of encoding the constituent locations, which typically would not affect processing if an intact ventral system were used.

*Subjects.* The patient, R.V., was a right-handed, bilingual male who had worked in technical training at a large computer company. He had earned a Bachelor's degree and was working toward a Master's degree. At the age of 38 he suffered a left frontal infarct, which will be discussed in detail shortly. Six months later he presented with mild anomia and slight deficiencies in speech production. Caplan (1990) tested him extensively on the Caplan-Bub Aphasia Battery, and found that he had moderate comprehension difficulties as well. In addition, he failed to name 13% of simple line drawings of common objects in their picture naming task;

virtually all of these errors were on pictures of animals. He was 39 years old at the time of testing.

A CT scan revealed that R.V.'s lesion was focused in the left frontal lobe. The damaged area appeared to be a cone whose base rested on the head of the caudate nucleus and whose tip just touched cortex near the region of the third convolution of the frontal lobe. MRI allowed much greater precision in characterizing this focal lesion. The lesion included zones of frank cavitation as well as zones of T1 and T2 signal prolongation consistent with gliosis; it was centered in the centro-sylvian region of the left hemisphere. Its extent was maximum in the region of the frontal operculum where the zone of signal change and cavitation spanned the full thickness of the cerebral wall. At the cerebral surface, the lesion destroyed much of the inferior opercular sections of Brodmann (1909; Bailey and von Bonin, 1951) areas 46, 45, 6, 4, plus the superior extent of 43 within the insula under the rostral parietal operculum. It intruded minimally into fields 3, 1, 2, and 40 within the sylvian fissure. Subcortically, the entire caudate and lenticular nuclei rostral to the thalamus as well as the adjacent segment of the horizontal limb of the diagonal band of Broca were destroyed and replaced by cavitation. The intervening corona radiata, external sagittal stium and anterior limb of the internal capsule were either marked by signal change or also frankly cavitated. Involvement of these central white matter systems extended forward through the forceps major beyond the callosal commissure. Caudally, the lesion also destroyed much of the posterior limb of the internal capsule to the level of the pulvinar.

We also tested a group of 8 control subjects. These subjects were right-handed men who responded to advertisements posted in various locations around Harvard University. They were approximately R.V.'s age (average age 36.6 years, range 33 to 42), and each was either working towards a Bachelor's degree or had no more than a Master's degree. Comparing R.V.'s results to those from relatively few control subjects will produce conservative estimates of R.V.'s deficits, which is reasonable given the large number of experiments that we must conduct to converge on possible accounts for deficits.

*Experiment 1.* In this experiment, the subjects saw two shapes in sequence and decided whether they were the same or different. The shapes were formed by filling in selected cells in 4x5 grids. We reasoned that the presence of the grid lines would define distinct locations, making it easier for the dorsal system to encode the shapes as sets of filled cells. In contrast, we expected the presence of grid lines to tax the ventral system by providing spurious input to the object properties system, which would be slowed down by the necessary additional processing.

We expected different patterns of response times if the response was produced on the basis of location information than if it was produced on the basis of shape representations. Specifically, the shapes varied in complexity, as defined by the number of perceptual units, or sets of contiguous filled-in cells, that comprised them. If the subjects looked sequentially at each cell to encode a shape, then they would take more time to respond to more complex shapes. On the other hand, if they encoded a shape as a single unit, then we would not see such an effect of complexity.

And in fact, R.V. required linearly increasing amounts of time to evaluate the stimuli with different numbers of perceptual units, whereas the control subjects did not. This and the following tasks were sufficiently easy, however, that there were very few errors. We designed relatively easy tasks intentionally because we wanted



to observe differences in the factors that affect correct response times—and could not use times from trials on which errors were made.

Because the ventral and dorsal systems operate in parallel, the subject's performance will reflect properties of whichever system produces useful output first. One account for these results is that the grid lines allowed R.V.'s dorsal system to produce a representation of the locations of filled cells more quickly than his ventral system could encode the overall shape; the hypometabolism in the ventral system may have impaired his encoding representations of shapes when the grid lines were present. Thus, it is possible that R.V. evaluated the shapes as sets of filled cell locations, using the dorsal system. Indeed, when interviewed later, R.V. claimed that he had compared the patterns by noting which individual cells were filled.

If so, then the output from the dorsal system would actually underlie R.V.'s responses in this experiment. If R.V.'s decisions were based on such encodings, then his response times would be sensitive to variables that affect the ease of encoding locations, such as the number of locations that had to be registered.

This is only one of many possible accounts for the results, however. In particular, it is possible that the more complex stimuli simply taxed his sluggish object-properties encoding processes, and so required more time. This hypothesis was ruled out by the results of the following experiment.

*Experiment 2.* In this experiment we presented the same shapes, but removed the internal grid lines and all of the perimeter of the grid except for small "brackets" defining the four corners. We reasoned that eliminating the grid lines would make it much more difficult to encode the locations of the filled regions. Thus, we expected these stimuli to be more easily encoded by the ventral system.

By including or deleting grid lines, then, we tried to manipulate the speed of two processes that operate in parallel—one that encodes sets of locations and one that encodes representations of shapes. The brain presumably does not "choose" which type of information to use in making a judgment; rather, whichever type is available first will be used. Thus, we hoped to vary the relative speed of the two methods, which in turn would affect which sort of information was ultimately used to produce the response.

We now found no hint of an effect of complexity on R.V.'s response times, in sharp contrast to the results of Experiment 1. The control subjects continued to show no difference in response times for increasingly complex stimuli. When the grid lines were removed, the shapes could not be easily encoded as sets of locations. Thus, they were more likely to be encoded via the ventral system. When this occurred, the units could be perceptually organized into a single form, and the response times were not sensitive to the number of filled locations.

However, there are other possible explanations for these findings, which must be considered empirically. In the following experiment we provide more direct evidence that the presence of the grid lines has an effect either beyond or different from simply overloading visual memory.

*Experiment 3.* The logic of Experiments 1 and 2 rested on the idea that the ventral system extracts features, and was taxed when more features were present. The existence of a "feature" detection process in the ventral system makes sense for a number of reasons. In particular, the visual system must be capable of recognizing objects when they are seen at unusual orientations. Lowe (1987a, b) built a computer vision system that solved this problem by in part by extracting "nonaccidental" properties from the input. These properties usually arise reliably from specific

aspects of a stimulus, even when the object is seen from an unusual vantage point; for example, parallel edges tend to project roughly parallel lines, symmetrical shapes tend to project symmetrical edges, edges that intersect project points of intersection, and so on (for a summary, see Biederman, 1987).

Biederman (1987) summarizes good behavioral evidence that humans use nonaccidental properties when identifying objects: When these properties are removed, objects are almost impossible to identify. In addition, Petersen, Posner, Schulman, and Raichle (1990) found that when humans were asked to view letter-like forms while being PET scanned, the lateral occipital-temporal cortex was activated. This cortex appears to be within the region of the hypometabolic area we found in R.V.'s brain. Thus, it is possible that his process that extracts edges and lines is not as efficient as it is in normal brains, and can be overly taxed by spurious lines.

However, one could argue that the added grid lines simply overloaded visual memory. That is, perhaps they did not result in a switch to a location-based response, but rather simply slowed down processing in the ventral system because there was more information to store about the first stimulus of the pair. If the previous results were caused at least in part by a defective encoding process, then we should find similar patterns even when memory is not required.

In this experiment, the subjects saw one of the stimuli used in Experiment 1 along with an X mark, and simply indicated whether the X fell on or off the shape. If R.V.'s deficit were a problem in remembering the first stimulus of a pair, then it should not be evident in this experiment. On the other hand, if his problem were a deficit in encoding visual features, then we would expect to find again an effect of complexity when grid lines are included.

We found that when the grid lines were present, R.V.'s response times did in fact increase linearly with the complexity of the shape. In contrast, there was no effect of complexity on the control subjects' response times. Clearly, the deficit observed in Experiment 1 was not due solely to impaired memory. Even when we eliminated the memory component of the task, a deficit was still evident.

Thus, we have evidence that the grid lines affect something other than visual memory. We speculate that they not only impair a feature encoding process, but also *promote* a location-encoding process, and that the effects of the two processes working together underlie the results. This inference is also consistent with the fact that PET scanning indicated hypometabolism in R.V.'s occipital-temporal area, which is where visual features are hypothesized to be encoded (e.g., for a review, see Chapters 3 and 5 of Kosslyn & Koenig, 1992), whereas his parietal system was unaffected. Furthermore, our findings recall reports by Luria (1980) of described patients with lesions in the occipital-temporal area who could not name pictures of objects when random line fragments were placed over them. Presumably, these lines overloaded a damaged feature extraction process, just as the grid lines affected our patient.

*Experiment 4.* We hypothesized that the results of the previous experiment indicated that R.V. used the locations in the grid to encode the shapes. This experiment was the same as Experiment 3, except the stimuli appeared in brackets instead of grids to impede this strategy.

And in fact, we found that removing the grid lines did indeed affect R.V.'s performance in this experiment. The linear effect of complexity was replaced, for some unknown reason, by a quadratic trend. (One- and three-unit stimuli were

processed faster than two-unit stimuli.) The control subjects showed no hint of an effect of complexity on response time.

*Experiment 5.* In Experiments 1 and 3, we assumed that responses would be based on processing in the dorsal system in part because the grid cells facilitate encoding locations. But the grid lines may also tax the visual encoding processes, and it is possible that the lines themselves are the root of the problem. This experiment was designed to investigate that possibility. It was identical to Experiment 3, except that the grid lines were removed and random line fragments were placed over the stimuli. These fragments were irregularly positioned, and sometimes intersected with one another or with the gray stimulus pattern. They did not form distinct cells, eliminating the option to encode the pattern as a set of locations in a grid. Thus, these stimuli forced the subjects to encode the patterns as shapes, and should have made the task relatively difficult if the lines overloaded the encoding processes, making the more complex patterns especially difficult to encode.

We found that simply adding noise fragments did not cause complexity to affect R.V.'s times or the control subjects' times. Thus, the presence of additional lines per se cannot be responsible for the previously observed increase in response time with increasingly complex patterns in grids.

*Experiment 6.* We have hypothesized that R.V. encodes shapes in grids as sets of locations. If so, then he should store such stimuli in this way, in addition to encoding them as locations in the first place. In this experiment we considered R.V.'s processing when he did not have to encode shapes at the time of test. Rather, he was allowed to study a shape in a grid, it was removed, and then an X mark appeared in the otherwise empty grid. R.V. was asked whether the X mark would have fallen on the figure if it were present. If he is in fact storing the patterns as sets of filled locations, then we should find increased time with more complex patterns even here.

Indeed, we now make the same prediction for R.V. and the control subjects. Because the task can only be performed by remembering which cells were filled, this experiment encourages the control subjects to "simulate" R.V.'s usual performance by attending to specific cells in the grid. Thus, it is of interest to discover whether this task leads the control subjects to pay attention to distinct locations, and hence to require more time when more locations comprised a shape.

And in fact we found that R.V.'s times increased with complexity, as did the control subjects' times. Thus, the requirement to remember the locations occupied by the figure apparently led both R.V. and the control subjects to encode the patterns as sets of filled locations. In sharp contrast to their results in all the previous experiments, the control subjects now required increasing amounts of time to evaluate increasingly complex patterns. Thus, R.V.'s strategy is not outside the range of normal behavior, but rather apparently is used by him in more situations than by normal subjects.

*Experiment 7.* It was possible that our previous findings with R.V. reflect a different kind of deficit; perhaps he simply cannot attend easily to relatively large regions of space. If so, perhaps grid lines were more distracting near the fringes of his attention. Because more complex stimuli tended to extend over a larger region, such a problem would produce all of the results described so far: R.V. would have more difficulty with complex (i.e., large) stimuli when grid lines were present. This experiment was designed to discover whether R.V.'s "attention window" (Kosslyn

et al., 1990) had an abnormally restricted scope. The subjects studied four gray blocks that were positioned along the circumference of an invisible circle, which was relatively large or small. On half the trials, an "X" mark appeared in two blocks on opposite sides of the circle, and on the other half of the trials only one X mark appeared. The subjects responded "yes" if both Xs were present, and "no" if only one was present. If R.V.'s attention window size were restricted, such that it could not easily be enlarged to cover the entire area of the larger circle, then we should find impaired performance on trials with stimuli placed on the circumference of the larger circle.

However, when we compared R.V.'s responses for large and small arrays, we found no effect of size. The same results were found with the control subjects. Thus, R.V.'s earlier impaired performance cannot be ascribed, even in part, to a deficit in his ability to attend to larger regions of space. R.V.'s problem apparently arises when he can easily encode a pattern as a set of locations, a strategy that requires more time for more complex patterns.

In summary, we found that R.V. became progressively slower when more cells of a grid were filled by a pattern, whereas control subjects generally did not require more time with more complex patterns in grids. However, the control subjects did show such an effect in one case, when it was useful to pay attention to the locations occupied by a figure in order to perform the task. The findings taken together are most consistent with the hypothesis that the grid lines promoted encoding the shapes as sets of filled locations, and R.V. had impaired shape-encoding processing in his ventral system—as we expected given the location of the hypometabolism—and thus tended to encode the shapes as sets of locations in grids.

At first glance, it may seem surprising that hypometabolism in one cerebral hemisphere had any noticeable effects. However, we must note that the damage was to the *left* side. Posner, Petersen, Fox, and Raichle (1988) report more activity in the left occipital temporal area when subjects were PET scanned while they saw words; it is possible that this area has been "tuned" to encode lines and angles, some of the constituents of letters, during reading, and hence performance was impaired when grid lines were included. Anecdotally, it may be worth noting that although R.V. could read, he was very slow and awkward; prior to the stroke, he was an avid and rapid reader.

The present approach is a departure from the usual technique in neuropsychology of establishing pairs of dissociations and associations following brain damage (e.g., Caramazza, 1986). We recognize that lesions are often relatively large, and sometimes have remote effects by de-energizing other parts of the brain. In this research we found evidence of a system of functional impairments that appears to reflect dysfunction in the occipital-temporal junction area (which putatively implements a feature encoding process). This approach is admittedly more complex than the usual method in neuropsychology, but seems appropriate in an investigation of a marvelously complex organ, the brain.

#### *Patient E.B.: Effects of reduced activation level*

In this study we examined patient E.B., who appeared to be in the early phases of Alzheimer's disease. This patient was 70 years old at the time of testing, and a very successful architect. His premorbid IQ was approximately 140. An MRI scan revealed cortical atrophy and a PET scan revealed marked decreased blood flow in the brain. However, there is no evidence of focal damage. Thus, we studied E.B. to

examine possible effects of "reduced activation levels." Indeed, E.B. complains that he has generally slowed down. Although we are just now analyzing these data in detail, enough analyses have been completed to allow us to summarize the results.

Using the tasks administered to R.V., we showed that E.B.'s cognitive deficit is not consistent across all tasks; rather, he performed abnormally on some tasks and normally on others. Subsequently, Dr. Shirley Wray prescribed an amphetamine-based drug for the patient, and we retested him after he began taking it. His deficits disappeared. Thus, we have good grounds for inferring that the deficits were in fact due to something like "activation level." The important finding is apparently qualitative variations in behavioral deficits—impairments only on selected tasks—can be related to a quantitatively varying underlying factor rather than merely the presence or absence of discrete processing components.

Thus, our series of case studies is illustrating consequences of different types of underlying deficits. In addition, these data are helping us to elaborate our theory of high-level vision (Kosslyn, Flynn, Amsterdam & Wang, 1990; Kosslyn, 1991), and are providing support for critical aspects of the theory.

#### *Evidence for Processing Subsystems*

The grant proposal described a battery of tasks, each of which was designed to tap a specific aspect of processing. The Kosslyn et al. (1990) theory was used to guide task construction; our aim was to design a task for each major aspect of processing. For example, the attention task used to test R.V. (Experiment 7) was designed to investigate one property of the "attention window." Some of the tasks were used to test the patients described above.

Altogether, we developed 27 tasks, each of which was thought to tap a distinct aspect of information processing. Four of the tasks (two image generation and two attention-related) were analyzed in two ways, yielding a total of 31 dependent measures. Briefly, the tasks selectively stressed processing in all of the component subsystems proposed by Kosslyn et al. (1990; see also Kosslyn, 1991), including the "attention window" within the visual buffer, the subsystems of the dorsal and ventral visual systems, top-down processes used in object identification and imagery, and image transformation processes used in such tasks as scanning and rotation.

Each patient also participated in a simple response-key training program at the beginning of testing, in which he or she simply pressed one key when "YES" appeared on the screen and another when "NO" appeared; all patients we tested could do this task almost perfectly. We used these response times as a baseline, to control for differences in simple sensory-motor abilities. Thus, for each task we subtracted the response baseline from the mean response time for that task. (We are also considering other ways of analyzing this data, including using ratios instead of differences.)

These tasks were administered to 17 patients, some of whom had focal damage and some of whom had more general, degenerative diseases. In addition to G.A., R.V., and E.B., who are described in detail above, these patients were (all are right-handed unless noted otherwise):

KS—50, male, Doctoral degree; NL—52, female, Master's degree; MM—81, female, Master's degree; MH—82, female, high school graduate; SR—22, male, some college; PF—81, male, some high school (ambidextrous); BB—82, male, Bachelor's degree; JC—73, male, some college; KG—74, male; JF—65, female (converted left-

hander); HC—56, female, some high school; MB—81, female, some high school; LH—40, male, Master's degree; AC—59, female, high school graduate. Some of these patients did not receive all of the tasks.

In neuropsychology the existence of a *double dissociation* is often taken as strong evidence that two tasks rely on at least some distinct underlying processes. For example, if a patient with an occipital lobe lesion could do a language task but not a visual task, and a patient with a frontal lesion could do the same visual task but not the same language task, we could infer that the two tasks do not completely share processing components—if they did, then both patients should have been impaired on both tasks. (A "single dissociation" would simply reflect a single patient's ability to do one task but not another.)

Here, we defined a double dissociation as the existence, for a pair of tasks A and B, of a pair of patients 1 and 2, such that patient 1 did better on task A than on task B, and patient 2 did better on task B than on task A. Graphically, with patient on the x-axis, this could be expressed as an interaction with two lines, one for each task, that rise in opposite directions. In addition to the interaction, there can be a main effect of patient (one being faster or more accurate than the other both tasks), or of task (one being harder than the other for both patients), or both.

We began the analysis by converting the baseline-adjusted patient scores into z-scores separately for each task; thus, each score was replaced by its number of standard deviations away from the mean of the entire group of patients for that task. For each of our 465 pairs of measures ( $31 \times 30 / 2$ ), we have computed three different summary statistics thus far, as follows. (All calculations were done by custom-written software.)

First, we simply counted the number of pairs of patients who showed a double dissociation by the above definition.

Second, we defined a new statistic,  $k$ , as follows. We created a scatter plot of the patient z-scores and divided it into two groups with a standard 45 degree line passing through the origin. If the z-score on task A is the x-axis and the z-score on task B is the y-axis, then the points above this line represent the patients who did better on task B than on task A, and vice-versa for the points below the line. Next, treating the above and below sets of points as two separate groups, we computed the variance between the groups (using the unweighted mean) and the variance within the groups. We then set  $k$  to be the ratio of between-group to within-group variance. (Note that this is *not* an  $F$  ratio, as found in the analysis of variance, for several reasons, one being that the within variance is not a proper error term.)

$k$  is always positive, so the  $k$  value for a task pair indicates how strongly the two tasks dissociate from each other (in this group of patients). For example, if most of the patients tended to score equally, in terms of z-scores, on two tasks, then the scatterplot would show a single narrow elliptical cloud of points whose major axis was approximately parallel to the 45 degree line. The between-group variance in the numerator of  $k$  would be virtually zero, so  $k$  would be close to zero itself. By contrast, if about half of the patients did much better on one task than another, and vice-versa for the rest of the patients, the scatterplot would show one distinct cloud in the upper left and another in the lower right, the between variance would be high, and  $k$  would be large. An intermediate case would have half the patients doing better on one task than another, but the other patients doing about equally well on the two tasks. The between variance here would not be as large as in the previous situation, since one cloud will lie near the center, along the dividing line.

Third, in order to cull the most interesting examples from this large data set for further study, we noted which pair of patients had the strongest double dissociation. To do this we calculated the "strength" of each double dissociation as follows. For each patient who did both tasks in the pair, we subtracted the z-scores for the tasks, yielding a z-difference-score. The strength of the double dissociation was simply the difference between these z-difference-scores for each pair of patients. Thus, if two patients scored similarly on two tasks, the strength of that double dissociation should be near zero.

Next we addressed the problem of testing the significance of the first two statistics described above. We designed the following Monte Carlo simulation, which is currently being executed:

Working from the data set *after* computing the within-task z-scores, we randomly reassigned the scores to different tasks for each patient. For example, if patient 1 completed only tasks A, B, and C, we might assign his score on task A to task B, task B to task C, and task C to task A. (Computing z-scores first eliminates effects of task—for example, the image rotation tasks in the battery virtually always result in longer response times than the spatial relations encoding tasks—and randomizing within patients preserves the effects of patient.)

After doing this for all 31 dependent measures, we ran the analyses described above and saved the results. We plan to repeat this procedure 10,000 times in order to collect stable data about the likely distribution of  $k$  and the number of double dissociations. For each possible number of patients doing both tasks in a pair (which ranges from 2 to 17 in this case, and can be thought of as analogous to degrees of freedom) we will have a distribution. By identifying the upper and lower tails we will find, for example, what values of  $k$  are necessary for "significance" at the .05, .01, or other levels, and we can then re-examine the original, nonrandomized data to see which pairs of tasks yielded  $k$  values outside the range we would expect due to chance. (That is, more or less than would be expected due to random variations if there were no double dissociation in the population—if the tasks did not rely on distinct processing components. For example, if there is no population double dissociation involving a pair of tasks, it should not matter how patient scores are "assigned" to those tasks. Thus, we randomized scores among tasks within patients.) Similar inferences can be made with respect to the data on the number of double dissociations observed per task pair.

The preliminary results indicate that there are double dissociations among 89% of the tasks. We plan to complete the analyses described above and then repeat it using error rates instead of response times as the dependent measure (in subtraction as well as ratio). We will also explore using other dependent measures, including raw response times and error rates and difference scores between within-task difficulty levels. Additionally, we will perform multidimensional scaling on correlation matrices of the various performance measures.

#### *Computer Simulation Models*

Our guiding assumption is that vision, like all other complex mental functions, is accomplished by a "divide and conquer" approach: Processing is split into many small parts, each of which is relatively simple (cf. Maunsell & Newsome, 1988; Van Essen, 1985). As noted earlier, a major division of labor is accomplished by systems in the temporal and parietal lobes, with the former encoding object

properties (such as shape and color) and the latter encoding spatial properties (such as location and size; see Kosslyn, Flynn, Amsterdam, & Wang, 1990; Maunsell & Newsome, 1988; Ungerleider & Mishkin, 1982). Each of these major systems can be further divided into component subsystems. In the present research, we focus on the system that encodes spatial properties, and argue that this system is divided into at least two subsystems which compute different kinds of representations of spatial relations.

The research reported here builds on the analyses and findings of Kosslyn (1987) and Kosslyn, Koenig, Barrett, Cave, Tang, and Gabrieli (1989). We further develop their conception of the two kinds of spatial relations representations by considering implications of new results reported by Sergent (1991), additional analyses of what is required to build a system that behaves in particular ways, and additional facts about the neural substrate.

Computational considerations suggest that different kinds of representations of spatial relations are useful for different purposes. Consider two contexts in which we use spatial information. First, we must use spatial information to guide actions, ranging from moving our eyes to reaching to navigating. And in fact, many cells in the posterior parietal lobes appear to have some role in movement control, either firing prior to or after a movement or registering the position of an effector (see Andersen, 1987; Hyvarinen, 1982). For guiding action, metric spatial information must be specified; simply knowing that a table is "next to" a wall will not help one to walk right up to it without bumping the edge.

Second, we often need to encode spatial relations to identify an object or scene. For this purpose, we need not represent metric information precisely; differences in the precise positions of two objects or parts often are not relevant for distinguishing them from other objects or parts. Rather, we assign spatial relations to a category, such as "connected to," "left of," or "above." For some purposes, it may also be useful to assign a spatial relation to a distance category, such as "one inch away," but this sort of category must be distinguished from the kind of analog encoding of metric distance that is necessary to guide action (e.g., see Osherson, Kosslyn, & Hollerbach, 1989).

These considerations lead us to the hypothesis that the brain represents spatial relations in two ways. First, *coordinate* representations specify precise spatial location in a way useful for guiding action. The units of these representations are not equivalence classes, but rather delineate the finest possible division of space. These representations do not correspond to particular movements, but rather specify spatial coordinates in a way that can be used to guide a variety of movements. Second, *categorical* representations assign a range of positions to an equivalence class (such as "connected to," "above," or "in front of"). For many objects, parts have the same categorical relations, no matter how the object contorts; thus, categorical spatial relations representations are a critical aspect of a robust representation of an object's shape (cf. Marr, 1982). For example, even though its position in space varies widely, a cat's paw remains "connected to" its foreleg when the cat is curled up asleep, running, or batting an insect.

We distinguish coordinate representations from those used in recognition in part because the spatial information used to guide action appears to be "encapsulated" (e.g., McLeod, McLaughlin, & Nimmo-Smith, 1985). The information used to guide action is not readily accessible to the systems used to categorize stimuli. For example, consider how easily we can throw a wad of paper



into a container, compared to how much difficulty we have estimating the distance of the container with equal precision.

In addition, neuropsychological findings have supported the claim that the two kinds of spatial representations are encoded by separate processing subsystems. Hellige and Michimata (1989), Koenig, Reiss, and Kosslyn (1990), and Kosslyn et al. (1989) all found that subjects encode metric spatial relations relatively faster when the stimuli are presented initially to the right cerebral hemisphere (i.e., in the left visual field), whereas they encode some categorical spatial relations relatively faster when stimuli are presented initially to the left cerebral hemisphere (i.e., in the right visual field) or equally well in both hemispheres. Although the left-hemisphere advantage is rarely significant in a single experiment, a trend towards a left-hemisphere advantage was present in six experiments that used stimuli presented at low contrast to adult subjects (Hellige & Michimata, 1989; Experiments 1, 2, 3, 4 of Kosslyn et al., 1989; Koenig et al., 1990; Experiment 4 of Sergent, in press), whereas a trend towards a right-hemisphere advantage was evident only once, in a less than 1 millisecond difference (Experiment 4 of Sergent, 1991). According to the binomial distribution (sign test), the probability that this pattern of results is due to chance is only .06.

More recently, however, Sergent (1991) reports that this dissociation occurs only when stimuli are relatively degraded, and she infers that this result does not reflect a distinction between two qualitatively different ways of representing spatial relations. Sergent assumes only that when the stimuli are degraded, the right hemisphere can more effectively encode precise position. She offers as one piece of evidence against the distinction the fact that more difficult discriminations (defined by relative distance) affect categorical judgments as well as metric judgments, and also reports several failures to find hemispheric differences in the two kinds of processing.

In Kosslyn, Chabris, Marsolek and Koenig (in press) we report computer simulations that support three assertions. (See the article for detailed information on the methods and results.) First, there is a qualitative distinction between categorical and coordinate spatial relations representations; second, more difficult discriminations affect the computation of both categorical and coordinate spatial relations representations; and third, the effects of stimulus quality can be accounted for easily by reference to a simple computational mechanism. However, Sergent's (1991) findings have led us to reconsider Kosslyn's (1987) original motivation for the distinction between categorical and coordinate spatial representations, which was based on the idea that the left hemisphere is specialized for language and the right is involved in navigation.

*Experiment 1. Split networks*

We have argued that categorical and coordinate spatial relations are qualitatively distinct, which leads us to expect them to be encoded by different subsystems. Kosslyn et al. (1990) assume that these subsystems correspond to separate neural networks, each of which maps an input (in this case, a representation of a pair of locations) to an output (in this case, a representation of a spatial relation). Accordingly, it is appropriate to use computer simulations of "neural networks," which establish mappings from sets of stimuli to correct responses, to investigate whether the two types of computations are in fact qualitatively distinct.

In this experiment we provide computational evidence that at least one categorical relation, above/below, is qualitatively distinct from the specification of metric location, which we claim is a critical component of coordinate representations. We simulated the bar-and-dot tasks developed by Hellige and Michimata (1989) and also used by Kosslyn et al. (1989, Experiment 3) and Koenig et al. (1990). In these tasks, subjects saw a short horizontal bar and a dot; they were asked either to determine whether the dot was above or below the bar (the categorical task) or was within a fixed distance—for example, 3 mm—from the bar (the coordinate task). In our simulations, the categorical task required the network to judge whether a single activated input unit (the dot) was "above" or "below" a landmark (the bar); the landmark consisted of two activated input units flanked by one inactivated input unit on each side. The coordinate task required the network to judge whether a single activated input unit (the dot) was within or was outside four units of the landmark (the bar). We considered this a coordinate task because the network had to encode the finest possible distinctions among locations; in contrast, the above/below task required grouping the locations into categories.

We studied these tasks using the "partition" paradigm developed by Rueckl, Cave, and Kosslyn (1989), which compares the performance of two types of networks on a pair of tasks; each type performs both tasks simultaneously, using separate sets of output units. One type of network has all of the hidden units connected to all of the output units, whereas the other type has its hidden units split into two groups. In a split network, one group of hidden units is connected exclusively to the output units for one task, whereas the other group is connected exclusively to the output units for the other task. Consequently, the representations developed by the hidden units for one input/output mapping cannot be used for the other, and vice versa. In the unsplit networks, in contrast, the hidden units form a single homogeneous group that is fully connected to all the output units for both tasks. (Note that a split network with  $a$  hidden units in one partition and  $b$  hidden units in the other partition is mathematically identical to two separate networks, one with  $a$  and one with  $b$  hidden units.)

Following the reasoning of Rueckl et al. (1989), we expect that if two tasks rely on distinct computations, a split network should perform the necessary mapping better than should an unsplit network. The segregation of processing prevents patterns of weights that are useful for accomplishing one input/output mapping from interfering with those that are useful for accomplishing the other mapping. However, this effect may not be evident until the networks have enough hidden units, because a split network has the inherent disadvantage of having fewer connections (and consequently fewer weight space dimensions) than the corresponding unsplit network. With sufficient hidden units, the advantage of

separating distinct representations should overcome the disadvantage of fewer resources. Accordingly, we systematically vary the number of hidden units to ensure that our results generalize over a range of models.

To establish the input/output mappings, we used the backward error propagation algorithm of Rumelhart, Hinton, and Williams (1986), as modified by Stornetta and Huberman (1987). This algorithm is sometimes characterized as a "learning" procedure, and its behavior is often compared to that of biological systems that learn (e.g., Rumelhart & McClelland, 1986). However, we do not assume that the kind of learning performed by the networks has a direct relation to learning in actual neural networks in the brain. Instead, we use the difficulty of "learning" in the models solely as a measure of how difficult it is to establish a specific input/output mapping. We treat the amount of error after a fixed number of training trials as a measure of the difficulty of establishing the mapping (cf. Rueckl & Kosslyn, in press).

In Part 1 of this experiment, we compare the ease of establishing categorical and metric mappings in split and unsplit networks with various numbers of hidden units. Each network was trained to establish both types of mappings, and we observed the amount of error after a fixed number of trials. In Part 2, we consider the possibility that any advantage shown by split networks might have nothing to do with the distinct types of representations: Perhaps dividing resources between two tasks is always beneficial, regardless of the degree of similarity between the mappings. To address this question, we compared split and unsplit networks that performed two different variants of the metric task, using the same input patterns and network architectures as in Part 1.

The results support the hypothesis that categorical and coordinate spatial relations are computed by distinct processes. In Part 1 we found that networks in which the two types of representations are segregated perform both mappings better than otherwise equivalent networks in which the representations are intermingled in the same set of processing units and connection weights. In Part 2 we found that unsplit networks were superior when the tasks were similar, and so the result of Part 1 cannot be merely a consequence of a general advantage for segregated processing, but is related to the distinct nature of the tasks.

For the split networks in Part 1, the optimal ratio of dividing the hidden units was usually close to an even split between the two mappings. This finding may have been an accident of the particular way we set up the tasks, however (e.g., the size of the input array, the amount that the bar moved, and so on), and we do not wish to draw strong inferences from this result; indeed, Kosslyn et al. (1989) manipulated the difficulty of discrimination to make the two kinds of tasks equivalent or either one more difficult than the other. However, it is fortunate that the two mappings were roughly equally difficult in these networks because it suggests that our results do not somehow reflect the effect of having a difficult task mixed with an easier one.

#### *Experiment 2. Effects of "difficulty"*

Sergent (1991) assumed that processes that encode categorical spatial relations would not be influenced by the distance separating the two objects. However, many categorical spatial relations rely on dividing space into discrete bins, and this process may be more difficult when the boundaries of these regions must be delineated more precisely. For example, a dot can be classified as "above" or "below" a landmark by observing whether it falls into one of two pockets of space; although

any location within each bin is treated as equivalent, it may be more difficult to assign a dot to a category if the bins must be delineated carefully. Furthermore, even after the regions of space are delineated, it may be more difficult to assign a dot to a category if it falls near the boundary. Such effects are found in a wide range of categorization tasks (e.g., Smith & Medin, 1983).

In this experiment we test Sergent's assumption directly, comparing the effect of discriminability in individual networks that perform either the categorical or coordinate mapping (not both simultaneously, as in Experiment 1). We test these specialized networks on two complementary subsets of the complete categorical and coordinate tasks: those stimulus patterns that Sergent's results suggest should be relatively easy and those that her results suggest should be relatively difficult.

In the categorical task, the easy discriminations are those in which the dot is far from the bar, and the difficult discriminations are those in which the dot is near the bar. In the coordinate task, the easy discriminations are those in which the dot is far from the criterion distance (from the bar), and difficult judgments are those in which the dot is close to the criterion distance. This experiment allows us to address Sergent's intuitions directly.

The networks revealed that subtle discriminations impair both kinds of judgments. This finding is important because, intuitively, it is not clear that a categorical spatial relation should be established less easily for stimuli that fall in a smaller range of positions. The findings of the following experiments will lead us to characterize this sort of categorical spatial relation as delineating regions of space, which is consistent with the present findings; the more difficult it is to delineate the regions to be related, the more difficult it is to establish the mapping. Thus, Sergent's finding that humans display a similar effect does not undermine the distinction between categorical and coordinate spatial relations representations.

#### *Experiment 3. The possible role of receptive field size*

The distinction between categorical and coordinate spatial relations was formulated following an analysis of the purposes of different kinds of spatial representations. We hypothesized that coordinate representations play a special role in action control, whereas categorical representations play a special role in recognition and identification. In this experiment we consider some implications of these ideas in more detail.

Action control depends on precise representation of spatial location. One way to represent spatial location precisely depends on overlap among rather coarse representations of location (Hinton, McClelland, & Rumelhart, 1986). This sort of *coarse coding* underlies color vision, for example, where the three types of retinal cone cells have overlapping distributions of sensitivity to different wavelengths; it is this overlap that allows the three types of cones to encode a much larger range of colors. O'Reilly, Kosslyn, Marsolek, and Chabris (1990) used back propagation models to show that this mechanism was also an effective way to encode metric spatial location.

It is possible that differences in the use of coarse coding can account for the hemispheric differences in computing the two types of spatial relations. On the one hand, the right hemisphere may receive input from low-level visual neurons that have relatively large receptive fields (i.e., receive input from relatively large regions of space), which have a large degree of overlap. These broadly-tuned receptive fields would allow effective coarse coding, and might explain the right hemisphere's superior ability to encode precise location. On the other hand, the left hemisphere

may receive input from low-level visual neurons that have relatively small receptive fields, which do not overlap very much. Sets of these receptive fields would define particular areas, which could be used to specify regions that are above or below a reference point, left or right of a reference point, and so on. In the limit, if the receptive fields did not overlap at all and the categories corresponded to discrete regions of space, such mappings would be "linearly separable"—so straightforward that they can be accomplished by direct connections from the input units to the output units, without a hidden layer (see Rumelhart & McClelland, 1986).

The idea that differences in receptive field properties may be at the root of differences in hemispheric specialization for spatial encoding is intriguing for a number of reasons. First, it fits nicely with Sargent's (1987) finding that the left hemisphere encodes smaller, high-spatial frequency patterns better than the right, and the right hemisphere encodes larger, low spatial frequency patterns better than the left (see also Van Kleeck, 1989). Presumably, higher spatial frequencies are encoded by smaller receptive fields, and lower spatial frequencies are encoded by larger receptive fields.

Second, neurons with large receptive fields would be useful in preattentive processing; preattentive processes, almost by definition, must monitor a wide range of visual angle. Preattentive processing plays a critical role in controlling actions; one often moves one's eyes, head, and body towards a sudden movement or other change in stimulus qualities. Furthermore, one wants to look at or reach to an object with reasonably good accuracy, even if it is seen out of the corner of one's eye. Livingstone (personal communication) has suggested that the magnocellular ganglia may project preferentially to the right hemisphere (see Livingstone & Hubel, 1987); these neurons have relatively large receptive fields, and are probably involved in preattentive processing.

Thus, it is of interest that overlapping large receptive fields not only will allow the system to monitor a large area, but also can produce the necessary precision to guide an initial movement—even if a target is seen out of the corner of one's eye. In keeping with this idea, Fisk and Goodale (1988; see also Goodale, 1988) report that right-hemisphere-damaged patients have a deficit in the initial phases of reaching towards a visual target.

Kosslyn, Hillger, and Livingstone (in preparation) recently tested a prediction of the idea that the right hemisphere receives proportionally more input from neurons that have relatively large receptive fields than does the left hemisphere. They showed subjects pairs of small oriented line segments, and asked them simply to decide whether the two segments had the same orientation. The members of a pair were presented in succession, always lateralized to the same visual field. On half the trials, the lines in a pair were presented in the same or nearby locations (no more than 1° of visual angle apart); on the other half of the trials, the lines were presented in relatively distant locations (from 3 to 8° of visual angle apart).

Kosslyn et al. found that subjects evaluated nearby lines more quickly if they were shown initially to the left hemisphere (i.e., presented in the right visual field) than to the right, but evaluated distant lines more quickly if they were presented initially to the right hemisphere. Presumably, in the left hemisphere more of the input is received from low-level neurons that have relatively small receptive fields, which have relatively high average resolution per minute of arc (because less area is averaged over). Hence, this input would allow better comparisons of nearby segments. But if segments are too far apart, they will be encoded by different sets of

low-level neurons, and hence require more processing in higher areas. In contrast, if the right hemisphere receives input from neurons that have larger receptive fields, the same neurons would encode both stimuli even when they are relatively far apart. Although the large receptive fields would thus facilitate making comparisons when stimuli are distant, the resolution within these receptive fields is relatively low—and hence the right hemisphere will not compare stimuli in nearby locations as well as the left.

These results must be evaluated in the context of Kitterle, Christman, and Hellige's (1990) failure to find any difference in the sensitivities of the two hemispheres to different spatial frequencies in a simple detection task. The orientation comparison task requires memory, and clearly draws on higher visual areas involved in comparison. And in fact, Kitterle et al. found that subjects *identified* high spatial frequency gratings better when they were presented in the right visual field (and hence were seen initially by the left hemisphere), whereas in some conditions subjects identified low-spatial-frequency gratings better when they were presented in the left visual field (and hence were seen initially by the right hemisphere). Thus, consistent with Sergent's (1982) ideas, the hemisphere differences cannot be ascribed to low-level processing, but rather depend on high-level encoding and comparison processes.

In short, we hypothesized that the right hemisphere receives relatively more input than the left from low-level visual neurons that have large receptive fields. These large receptive fields overlap, which allows the right hemisphere to encode coordinate spatial relations better than the left. In contrast, the left hemisphere receives relatively more input from low-level neurons with small receptive fields. These receptive fields do not overlap as much as the larger ones, which allows the left hemisphere to specify some categorical relations by delineating discrete sets of locations. For example, if an X is "left of" a Y, in many cases one may be able to represent the relation by defining two regions, one for the left and one for the right.

On this view, if a categorical relation cannot be computed by defining discrete pockets of space, then the left hemisphere will not encode categorical relations better than the right. This hypothesis is consistent with Sergent's (1991) failure to find left-hemisphere superiority in tasks that used spatially complex stimuli, which will not allow easy delineation of specific regions of space.

These hypotheses rely on subtle distinctions and several steps of reasoning. Thus, they are ideal candidates for computer simulation modeling; such modeling will show whether they are so much hand-waving or whether they have any clear implications. In this experiment we used network models to test these hypotheses in two ways. In Part 1, we examined the mappings performed by the categorical and coordinate networks by analyzing the "receptive fields" developed by the hidden units of different networks. That is, for each hidden unit, we examine which regions of the input array most strongly influence its level of activation. The stronger the weight on the connection from an input unit to the hidden unit, the more strongly a dot in that location will affect the hidden unit; thus, the pattern of weights on the connections to a hidden unit define its "receptive field" (cf. Lehky & Sejnowski, 1988; Zipser & Andersen, 1988). Specifically, we test the possibility that coordinate networks develop relatively larger receptive fields than categorical ones, after training is allowed to proceed until the networks had zero error.

As predicted, the networks that performed the categorical task spontaneously developed significantly larger receptive fields than the networks that performed the

metric task. However, as was found by O'Reilly et al. (1990), these receptive fields often tended to have complex shapes. It is possible that the size differences are somehow related to the various sets of shapes that developed. To examine the effects of size per se, we manipulate only this variable in Part 2 of this experiment.

In Part 2 we construct networks that are "hard-wired" to have relatively large or small receptive fields, and consider how effectively they perform the two kinds of mappings. Mirroring the results from Part 1 of this experiment, networks with fixed large receptive fields performed the metric task better than networks with fixed small receptive field sizes, and there was a tendency for the opposite pattern in the categorical task. These findings are like those from the corresponding experiments with human subjects. Indeed, only one out of the four individual experiments that used the above/below and near/far tasks showed a significant left-hemisphere advantage for the categorical judgment. However, three out of the four experiments revealed a trend towards a left hemisphere advantage, with the one deviant revealing a less than one millisecond difference. (In all four of those experiments, the interaction of task and hemisphere was significant; in these simulations the corresponding interaction of task and receptive field size was significant as well.)

Thus, these results provide support for the claim that distinct computations encode categorical and coordinate spatial relations representations. Furthermore, they hint at one possible reason why the hemispheres are specialized for encoding the different types of representations.

#### *Experiment 4. Accounting for the effect of contrast*

The results of the previous experiment may provide an insight into Sergent's (1991) finding that hemispheric differences do not arise when the stimuli have high contrast. We have suggested that the hemispheres differ in the input they receive from neurons that have different sized receptive fields. Kosslyn et al. (1989) suggested that the hemispheres differ in their *relative* efficacy at encoding the two types of spatial relations. (They did not claim that the hemispheres were exclusively specialized for the different types of encoding, as Sergent [1991] apparently inferred.) This leads us to expect that the difference in inputs to the two hemispheres is one of degree, with the left hemisphere receiving more input from neurons with relatively small receptive fields and the right hemisphere receiving more input from neurons with relatively large receptive fields. Hence, we expect a distribution of inputs from neurons with different receptive field sizes in both hemispheres.

These ideas suggest that the modulation transfer functions (see Kaufman, 1974) of high-level visual areas may differ in the two hemispheres. We are not considering the modulation transfer functions of low-level areas involved in detection, but rather those of higher areas involved in memory and comparison. The idea is that the modulation transfer functions of the high-level areas are slightly shifted, so that the peak sensitivity for the right hemisphere is at a lower spatial frequency, which reflects receiving more input from larger receptive fields. At relatively low contrast, the performance of the two hemispheres would be well separated; however, at relatively high contrast (when less sensitivity is required), the two distributions have a large amount of overlap—and hence the performance of the two hemispheres would not be well separated.

This could explain Sergent's (1991) finding that hemispheric differences between categorical and coordinate processing only occur when the stimuli are relatively degraded (i.e., when they are presented at low contrast). Hence, we decided to test this hypothesis with another set of simulated networks. In these models, we

again varied the sizes of the "receptive fields" of the inputs, and compared networks in which a narrow range of receptive field sizes are used (corresponding to low contrast) to those in which a relatively wide range of receptive field sizes are used (corresponding to high contrast).

In addition, we compared this hypothesis to a simpler one: Perhaps increased contrast does not recruit additional neurons that increase the range of receptive field sizes, but rather eliminates hemispheric differences simply because more of the same type of low-level neurons are stimulated over threshold. Even if neurons have relatively small receptive fields, enough overlapping outputs would allow coarse coding to be used effectively even in the left hemisphere.

Thus, we compared two sets of networks: ones in which greater contrast was assumed to produce outputs from a wider range of sizes of receptive fields ("mixed" receptive field size networks), and ones in which greater contrast was assumed simply to produce more input from additional neurons (with the same-sized receptive fields) in lower visual areas ("homogeneous" receptive field size networks). If the first hypothesis is correct, the effect of receptive field size found in Experiment 3 should be eliminated in mixed networks only; if the second hypothesis is correct, we should not find this effect in homogeneous networks.

As it turned out, we were best able to model the behavioral findings if we simply assumed that greater contrast leads simply to greater output from low-level neurons. Adding inputs to the network effectively eliminated the advantage of large receptive fields for encoding precise location; the networks were able to use coarse coding effectively, even if the receptive fields were relatively small. This observation makes sense when one considers that coarse coding is an effective strategy only when *sufficient* and *systematic* overlap exists in the distribution of response profiles. It has previously been demonstrated that overlap in receptive fields must be systematic in order to encode precise locations (O'Reilly et al., 1990; see also Ballard, 1986). Our findings demonstrate that the overlap must also reach a sufficient *amount* for encoding to be effective.

Presenting stimuli with higher levels of contrast probably results in a greater number of response profiles entering the distribution. Therefore, if the added receptive fields represent areas located fairly evenly across the input space, a higher and more effective degree of overlap is obtained even among relatively small receptive fields.

In summary, the present results lead to several general conclusions. First, we have additional support for the distinction between categorical and coordinate spatial relations encoding. Not only did we provide computational evidence that the two kinds of spatial relations are qualitatively distinct, but we also were able to formulate a simple mechanism that would explain why the lateralization of the two types of representations is not evident when stimuli are presented with high contrast. Our findings suggest that the distinction between categorical and coordinate spatial relations encoding is both principled and computationally plausible.

#### *Spatial Abilities in Pilots*

Finally, as a spin-off from our work with patients, we decided to test air force pilots on some of our visual-spatial abilities tasks. Drs. Alfred Fregly and Wayne Waag were instrumental in allowing us to have access to pilots. After graduate student Itiel Dror and I visited Williams AFB and discussed our plans, Dr. Waag



invited Mr. Dror to visit for an extended stay to test a group of pilots. This research appears quite promising.

#### *Subjects*

Sixteen pilots and 16 control subjects were tested in five experiments. The pilots and four of the control subjects were recruited and tested at the Aircrew Training Research Division of the Armstrong Laboratory in Williams Air Force Base, AZ. All 16 pilots were male, their mean age was 30 (range 23–46), 14 were right- and two were left-handed, and they all had at least a college education. The pilots had a mean of 1773 flight hours (range 218–4170, with only three having less than 1000). The control subjects were members of the Harvard University community and were tested on campus. The sex, age, handedness, and education of the control group was matched to those of the pilots: They were all male, their mean age was 29 (range 21–44), 14 were right- and two were left-handed, and they all had at least a college education.

#### *Experiment 1: Mental rotation*

Visual mental images can be transformed in many ways. One transformation that has received much attention is image rotation, which requires subjects to imagine an object rotating. Previous research has shown that subjects require progressively more time to image rotating an object progressively greater amounts (e.g., Shepard & Cooper, 1982). The time to alter the orientation is reflected by the slope of the increase in response times for greater rotation, and many researchers have compared the slopes of subjects in different groups (e.g., see Dror & Kosslyn, 1991; Kosslyn et al., 1990).

Our image rotation task required the subjects to determine whether two sequentially presented shapes were identical, regardless of their orientation. The shapes were rotated only in the picture plane; the first was upright and the second was at one four angular disparities relative to the first. This task is a modified version of the task devised by Shepard and Metzler (1971), which revealed a linear increase in response time with the increased angular disparity.

We found that pilots can rotate objects in mental images much better than can non-pilots. Indeed, about half of the pilots showed no increase in time to rotate objects greater amounts! Although the complete task probably recruits many processing subsystems, only a few are used to actually rotate the image. Thus, the difference in rotation slopes allows us to narrow down the pilots' advantage to those processes that change the represented orientation per se. Indeed, Kosslyn (1987) claims that only three subsystems are used to rotate imaged patterns, one that shifts the represented positions of segments, one that monitors the spatial relations among segments, and one that looks up stored information—which is used to direct the shifting subsystem to realign the segments properly. The results of additional tasks, described below, indicate that pilots are not better at encoding the kinds of spatial relations that would be used to note the relative positions of the segments ("categorical" relations, indicating that the segments are preserving right angles), nor are they better at accessing spatial information in memory. Thus, we can infer that the present results reflect a selectively superior ability to shift the locations of segments of the imaged pattern. We also found that pilots were faster overall than non-pilots—which could reflect in part their superior mental rotation ability, but also could reflect the effects of other subsystems that are used to encode stimuli or make responses.

#### *Experiment 2: Trajectory projection*

Another basic visual-spatial ability allows one to track a moving object, and to anticipate its trajectory when it is no longer visible. We administered a task that required the subjects to track a ball on the computer screen. The ball traveled at a constant speed for about two full circles, and then disappeared from the screen. After a delay, an X mark was presented briefly and the subjects decided whether it would have covered the ball if it had continued on its previous path at the same speed.

The motion task did not reveal any difference between pilots and non-pilots. Although increasing delays and distances of the X probe did affect their performance, they did so in the same way for subjects in both groups. However, this was a difficult task, as witnessed by the high error rates even in for "easy" trials. Thus it is possible that we would have found a difference if we had not taxed processing so much. The varying levels of difficulty between the trials might have not been large enough to allow us to detect differences between the groups (there was only 732 ms difference between the shortest and longest delay, and only 12 or 24 degrees of disparity between the X mark and the position of the ball). As usual, we must be cautious before affirming the null hypothesis, so further investigation of this task is indicated.

#### *Experiment 3: Scanning images*

In the previous experiment the subjects were asked to visualize the ball moving along its path after it was no longer actually on the screen. This process may rely on a specific subsystem that is used in visual mental imagery, namely that which allows one to scan over an imaged object. This sort of scanning appears to be used for a wide variety of purposes. Shepard and Cooper (1982) provide many examples of cases in which such visual mental imagery was used in scientific problem-solving, and it seems likely that imagery is used in the service of many kinds of reasoning (see Chapter 4 of Kosslyn and Koenig, 1992). This sort of reasoning depends on visualizing possible scenarios and "seeing" what would happen, "inspecting" the imaged patterns—and the process of image inspection often involves scanning over the image. Thus, we decided to assess image scanning ability in the two groups.

Our image scanning task was a variant of one developed by Finke and Pinker (1982). The subjects saw a rectangular-ring shape, which was composed of white and black squares. An arrow appeared briefly within the ring, after which the entire display disappeared, and the subjects decided whether the arrow had been pointing to a black square. We varied the distance from the arrow to the square ring, expecting more scanning when farther distances had to be "traversed." And in fact, Finke and Pinker's (1982) subjects required more time when greater distances had to be scanned (replicating the original finding of Kosslyn, 1973).

We found that pilots and non-pilots scanned images at comparable rates. This result is interesting in light of our findings in the motion tracking task, where increasing delays affected performance in the same way for subjects in both groups. Unlike that task, however, this one was not so difficult as to suggest possible floor effects. Thus, we are led to suspect that pilots and non-pilots can mentally extrapolate trajectories equally well.

#### *Experiment 4: Spatial relations encoding*

Given our findings (noted above and in Kosslyn et al., in press) supporting the claim that there are two ways to represent spatial relations, it was of interest to assess how well pilots and non-pilots could encode the two kinds of spatial relations. We administered two spatial relations tasks, one that required the subjects

to make categorical spatial judgments and another task that required them to make metric spatial judgments. We used the same tasks Kosslyn et al. (in press) simulated with their "neural network" models. The categorical task required the subjects to decide whether an X mark was above or below a bar; the distance of the X mark from the bar was varied to manipulate difficulty. The same stimuli were also used in the metric task, but now the subjects were required to decide whether the X mark was within half an inch of the bar; the distance of the X mark from this criterion point was varied to manipulate difficulty (see also Kosslyn et al., 1989).

We found that the pilots judged metric distance better than non-pilots, and again were generally faster than the non-pilots. Given the previous neuropsychological findings, the results suggest that the pilots may be better than non-pilots at specific kinds of right-hemisphere processing. Moreover, there is good evidence that spatial abilities depend critically on the parietal lobes (for reviews see Andersen, 1987; De Renzi, 1983; Ungerleider & Mishkin, 1982). Thus, it seems likely that pilots have superior coordinate spatial relations subsystems located in the right parietal lobe.

#### *Experiment 5: Recovering visual features*

We often see objects when they are partially occluded or obscured by various kinds of visual noise. For example, an object might be partially behind a bush, off in the distance on a foggy day, and so on. Lowe (1987), Biederman (1987), and others have proposed that the visual system copes with such problem by extracting relatively invariant features, such as parallel edges, points where edges intersect, and symmetries. We examined the ability to recover such features in the face of visual noise. In the "easy" condition, the subjects saw a curvy shape and an X mark, and were asked to decide whether the X mark was on the shape; in the "difficult" condition, the subjects performed the same task except that "noise" lines were imposed randomly over the stimulus. (See the description of patient R.V. above for more information on this type of task; see also Kosslyn, Daly, et al., in press.)

We found that pilots are no better than non-pilots at extracting visual features from images through random noise. This is of interest in part because different brain mechanisms are used to extract such features than are used to process spatial properties. Indeed, as discussed previously, Ungerleider and Mishkin (1982) and others document that "what" and "where" are processed by different anatomical pathways, with the former kind of information being the province of the temporal lobes and the latter the province of the parietal lobes. The fact that we have found that the pilots are better than nonpilots only on spatial tasks is intriguing given this division of labor.

Thus, we have shown the utility of the concepts and methodologies developed to study brain-damaged patients. These tasks show promise of allowing us to characterize what is special about the visual-spatial abilities of pilots, and perhaps other groups of professionals. After we have a more complete profile of the special visual-spatial strengths of pilots, the next step will be to discover whether these abilities are developed in the course of training, or characterize good candidates before they begin training.

#### *Concluding Remarks*

In short, we accomplished much in 1991. We continue to use an interplay of empirical work with brain-damaged patients and computational modeling to

develop our ideas about higher-level processing. We anticipate doing more of the same during the coming year.

## References

- Allman, J. M., and Kaas, J. H. (1976). Representation of the visual field on the medial wall of occipital-parietal cortex in the owl monkey. *Science*, 191, 572-575.
- Andersen, R. A. (1987). The role of the inferior parietal lobule in spatial perception and visual-motor integration. In F. Plum, V. B. Mountcastle, and S. R. Geiger (Eds.), *Handbook of Physiology: The Nervous System V, Higher Functions of the Brain Part 2*. Bethesda, MD: American Physiological Society.
- Andersen, R. A., Essick, G. K., and Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, 230, 456-458.
- Attneave, F. (1974). How do you know? *American Psychologist*, 29, 493-499.
- Bailey, P., and von Bonin, G. (1951). *The Isocortex of Man*. Urbana, IL: University of Illinois Press.
- Bauer, R. M., and Rubens, A. B. (1985). Agnosia. In K. M. Heilman and E. Valenstein (Eds.), *Clinical Neuropsychology*. New York: Oxford University Press.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Bower, G. H., and Glass, A. L. (1976). Structural units and the reintegrative power of picture fragments. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 456-466.
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde*. Leipzig: Barth.
- Campion, J. (1987). Apperceptive agnosia: the specification and description of constructs. In Humphreys, G. W., and Riddoch, M. J. (1987a) (Eds.), *Visual Object Processing: A Cognitive Neuropsychological Approach*. Hillsdale, NJ: Erlbaum. pp. 197 - 232.
- Caramazza, A. (in press). Is cognitive neuropsychology possible? *Journal of Cognitive Neuroscience*
- Cohen, G. (1975). Hemispheric differences in the utilization of advance information. In P. M. A. Rabbit and S. Dornic (Eds.), *Attention and Performance V*. New York: Academic Press.
- Cowey, A. (1985). Aspects of cortical organization related to selective attention and selective impairments of visual perception: a tutorial review. In M. S. Posner and O. S. Marin (Eds.), *Attention and Performance XI*. Hillsdale, NJ: Erlbaum.
- Critchley, M. (1953). *The Parietal Lobes*. Edward Arnold: London.
- Damasio, A. R. (1986). Disorders of complex visual processing: agnosias, achromatopsia, Balint's syndrome and related difficulties of orientation and construction. In M.-M. Mesulam (Ed.), *Principles of Behavioral Neurology*. Philadelphia, PA: F. A. Davis. (pp 259-288)
- Damasio, A. R., Damasio, H., and Van Hoesen, G. W. (1982). Prosopagnosia: Anatomic basis and behavioral mechanisms. *Neurology (NY)*, 32, 331-341.
- De Renzi, E. (1982). *Disorders of Space Exploration and Cognition*. New York: John Wiley & Sons.
- Deutsch, G., Bourbon, W. T., Papanicolaou, A. C., & Eisenberg, H. M. (1988). Visuospatial experiments compared via activation of regional cerebral blood flow. *Neuropsychologia*, 26, 445-452.
- Farah, M. J., and Hammond, K. M. (1988). Mental rotation and orientation-invariant object recognition: dissociable processes. *Cognition*, 29, 29 - 46.

- Fisk, J. D., and Goodale, M. A. (1988). The effects of unilateral brain damage on visually guided reaching: hemispheric differences in the nature of the deficit. *Experimental Brain Research*, 72, 425-435.
- Frackowiak, R.S., Lenzi, G.L., Jones, T., and Heather, J.D. (1980). Quantitative measurement of regional cerebral blood flow and oxygen metabolism in man using <sup>15</sup>O and positron emission tomography: Theory, procedure and normal values. *Journal of Computer Assisted Tomography*, 4, 727-726.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational knowledge. In F. Plum and V. Mountcastle (Eds.), *Handbook of Physiology*, Vol 5. Washington, D.C.: American Physiology Society. (pp. 373-417)
- Goodale, M. A. (1988). Hemispheric differences in motor control. *Behavioural Brain Research*, 30, 203-214.
- Gould, S. J., and Lewontin, R. C. (1978). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London*, 205, 581-598.
- Gregory, R. L. (1970). *The Intelligent Eye*. London: Weidenfeld and Nicholson.
- Hannay, H. J., Varney, N. R., and Benton, A. L. (1976). Visual localization in patients with unilateral brain disease. *Journal of Neurology, Neurosurgery and Psychiatry*, 39, 307-313.
- Hecaen, H., and Albert, M. L. (1978). *Human Neuropsychology*. New York: John Wiley and Sons.
- Hellige, J. B. and Michimata, C. (1989). Categorization versus distance: Hemispheric differences for processing spatial information. *Memory & Cognition*, 17, 770-776.
- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart and J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Hock, H., Kronseder, C., and Sissons, S. (1981). Hemispheric asymmetry: the effect of orientation on same different comparison. *Neuropsychologia*, 19, 723-727.
- Holmes, G. (1919). Disturbances of visual space perception. *British Medical Journal*, 2, 230-233.
- Humphreys, G. W., and Riddoch, M. J. (1987a) (Eds.). *Visual Object Processing: A Cognitive Neuropsychological Approach*. Hillsdale, NJ: Erlbaum.
- Humphreys, G. W., and Riddoch, M. J. (1987b). *To See But Not To See: A Case Study of Visual Agnosia*. Hillsdale, NJ: Erlbaum.
- Hyvarinen, J. (1982). Posterior parietal lobe of the primate brain. *Physiological Review*, 62, 1060-1129.
- Jackson, J. H. (1874). On the duality of the brain. Reprinted (1932) in J. Taylor (Ed.), *Selected Writings of John Hughlings Jackson*. London: Hodder and Stoughton.
- Kaufman, L. (1974). *Sight and Mind*. New York: Oxford University Press.
- Kinsbourne, M., and Warrington, E. K. (1963). The localizing significance of limited simultaneous visual form perception. *Brain*, 86, 697-702.
- Kitterle, F. L., Christman, S., and Hellige, J. B. (1990). Hemispheric differences are found in the identification, but not the detection, of low versus high spatial frequencies. *Perception and Psychophysics*, 48, 297-306.
- Kolb, B. and Whishaw, I. Q. (1985). *Fundamentals of Human Neuropsychology* (2nd Ed.). San Francisco: W. H. Freeman.

- Kolb, B. and Whishaw, I. Q. (1985). *Fundamentals of Human Neuropsychology* (2nd Ed.). San Francisco: W. H. Freeman.
- Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. *Psychological Review*, 94, 148-175.
- Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: a computational approach. *Psychological Review*, 94, 148-175.
- Kosslyn, S. M. (1988). Aspects of a cognitive neuroscience of mental imagery. *Science*, 240, 1621-1626.
- Kosslyn, S. M. (1991). A cognitive neuroscience of visual mental imagery: further developments. In R. Logie (Ed.), *Advances in Mental Imagery Research*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M. (in press). A cognitive neuroscience of visual mental imagery: further developments. In R. Logie (Ed.), *Advances in Mental Imagery Research*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M., Berndt, R. S., and Doyle, T. J. (1985). Imagery and language: a preliminary neuropsychological investigation. In M. S. Posner and O. S. Marin (Eds.), *Attention and Performance XI*. Hillsdale, NJ: Erlbaum.
- Kosslyn, S. M., Cave, C. B., Provost, D., and Von Gierke, S. (1988). Sequential processes in image generation. *Cognitive Psychology*, 20, 319-343.
- Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., and Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*, 34, 203-277..
- Kosslyn, S. M., Hillger, L. A., and Livingstone, M. (in preparation). Hemispheric differences in sizes of receptive fields? Harvard University manuscript.
- Kosslyn, S. M., and Intriligator, J. M. (in press). Is cognitive neuropsychology plausible? The perils of sitting on a one-legged stool. *Journal of Cognitive Neuroscience*.
- Kosslyn, S. M., and Koenig, O. (1992). *Wet Mind: The New Cognitive Neuroscience*. New York: Free Press.
- Kosslyn, S. M., Koenig, O., Barrett, A., Cave, C. B., Tang, J., and Gabrieli, J. D. E. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 723-735.
- Kosslyn, S. M., and Van Kleeck, M. (1990). Broken brains and normal minds: Why humpty-dumpty needs a skeleton. In E. Schwartz (Ed.), *Computational Neuroscience*. Cambridge, MA: MIT Press.
- Krieg, W. J. S. (1973). *Architectonics of the Human Cerebral Fiber Systems*. Evanston, IL: Brain Books.
- Larsen, A., and Bundesen, C. (1978). Size scaling in visual pattern recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 1 - 20.
- Lehky, S. R. and Sejnowski, T. J. (1988). Neural network model for cortical representation of surface curvature from images of shaded surfaces. In J. S. Lund (Ed.), *Sensory Processing in the Mammalian Brain*. Oxford: Oxford University Press.
- Levine, D. N. (1982). Visual agnosia in monkey and man. In D. J. Ingle, M.A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press.
- Levine, D. N., Maini, R. B., and Calvanio, R. (1988). Pure agraphia and Gerstmann's Syndrome as a visuospatial-language dissociation: An experimental case study. *Brain and Language*, 35, 172-196.

- Litton, J., and Bergstrom, M., and Eriksson, L., et al. (1984). Performance study of the PC-384 PET camera for the brain. *Journal of Computer Assisted Tomography*, 8, 74-87.
- Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science*, 240, 740-749.
- Lowe, D. G. (1987a). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31, 355-395.
- Lowe, D. G. (1987a). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31, 355-395.
- Lowe, D. G. (1987b). The viewpoint consistency constraint. *International Journal of Computer Vision*, 1, 57-72.
- Luria, A. R. (1959). Disorders of "simultaneous perception" in a case of bilateral occipitoparietal brain injury. *Brain*, 82, 437-449.
- Luria, A. R. (1980). *Higher Cortical Functions in Man*. New York: Basic Books.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- Maunsell, J. H. R. and Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, 10, 363-401.
- McLeod, P., McLaughlin, C., and Nimmo-Smith, I. (1985). Information encapsulation and automaticity: evidence from the visual control of finely timed actions. In M.I. Posner & O.S.M. Marin (Eds.), *Attention and Performance XI*, 391-406. Hillsdale, NJ: Erlbaum.
- Mehta, Z., Newcombe, F., and Damasio, H. (1987). A left hemisphere contribution to visuospatial processing. *Cortex*, 23, 447-461.
- Mesulam, M-M. (1981). A cortical network for directed attention and unilateral neglect. *Annals of Neurology*, 10, 309-325.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782-784.
- O'Reilly, R. C., Kosslyn, S. M., Marsolek, C. J., & Chabris, C. F. (1990). Receptive field characteristics that allow parietal lobe neurons to encode spatial properties of visual input: a computational analysis. *Journal of Cognitive Neuroscience*, 2, 141-155.
- Olson, D., and Bialystok, E. (1983). *Spatial Cognition: the Structure and Development of Mental Representations of Spatial Relations*. Hillsdale, NJ: Erlbaum.
- Osherson, D. N., Kosslyn, S. M., and Hollerbach, J. M. (Eds.) (1989), *An Invitation to Cognitive Science, Volume II: Visual Cognition and Action*. Cambridge, MA: MIT Press.
- Palmer, S. E. (1977). Hierarchical structure in perceptual representations. *Cognitive Psychology*, 9, 441-474.
- Pandya, D. N., and Yeterian, E. H. (1985). Architecture and connections of cortical association areas. In A. Peters & E. G. Jones (Eds.), *Cerebral Cortex, Vol 4: Association and Auditory Cortices*. New York: Plenum.
- Pohl, W. (1973). Dissociation of spatial discrimination deficits following frontal and parietal lesions in monkeys. *Journal of Comparative and Physiological Psychology*, 82, 227-239.
- Posner, M. I., Inhoff, A. W., Friedrich, F. J., and Cohen, A. (1987). Isolating attentional systems: A cognitive-anatomical analysis. *Psychobiology*, 15, 107 - 121.
- Posner, M. I., Petersen, S. E., Fox, P. T., and Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, 240, 1627-1631.



- Posner, M. I., Petersen, S. E., Fox, P. T., and Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science*, 240, 1627-1631.
- Posner, M. I., Walker, J. A., Friedrich, F. J., and Rafal, R. D. (1984). Effects of parietal lobe injury on covert orienting of visual attention. *Journal of Neuroscience*, 4, 1863-1974.
- Ratcliff, G. (1982). Disturbances of spatial orientation associated with cerebral lesions. In M. Potegal (Ed.), *Spatial Abilities: Developmental and Physiological Foundations*. New York: Academic Press.
- Ratcliff, G., and Davies-Jones, G. A. G. (1972). Defective visual localisation in focal brain wounds. *Brain*, 95, 49-60.
- Reed, S. K., and Johnsen, J. A. (1975). Detection of parts in patterns and images. *Memory and Cognition*, 3, 569-575.
- Rosenthal, R. (1984). *Meta-analytic Procedures for Social Research*. Beverly Hills, CA: Sage.
- Rueckl, J. G. and Kosslyn, S. M. (in press). What good is connectionist modelling? In A. Healy, R. M. Shiffrin, and S. M. Kosslyn (Eds.), *Essays in Honor of W. K. Estes*. Hillsdale, NJ: Erlbaum.
- Rueckl, J. G., Cave, K. R., and Kosslyn, S. M. (1989). Why are "what" and "where" processed by separate cortical visual systems? A computational investigation. *Journal of Cognitive Neuroscience*, 1, 171-186.
- Rumelhart, D. E., Hinton, G., and Williams, R. (1986). Learning internal representations by error propagation. In D. E. Rumelhart and J. L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Senda, M., Buxton, R. B., Alpert, N. M., and Correia, J. A., Mackay, B. C., Weise, S. B., and Ackerman, R. H. (1988). The <sup>15</sup>O steady-state method: Correction for variation in arterial concentration, *Journal of Cerebral Blood Flow and Metabolism*, 8, 681-690.
- Sergent, J. (1982). The cerebral balance of power: Confrontation or cooperation? *Journal of Experimental Psychology: Human Perception and Performance*, 8, 253-272.
- Sergent, J. (in press). Judgments of relative position and distance on representations of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance*.
- Shepard, R. N., and Cooper, L. A. (1982). *Mental Images and their Transformations*. Cambridge, MA: MIT Press.
- Stornetta, W. S. and Huberman, B. A. (1987). An improved three-layer back propagation algorithm. *Proceedings of the IEEE First International Conference on Neural Networks*, San Diego, CA.
- Talairach, J., and Sziklag, L. (1967). *Atlas of Stereotaxic Anatomy of the Telencephalon*. Paris: Masson and Cie.
- Taylor, A. M., and Warrington, E. K. (1973). Visual discrimination in patients with localized brain lesions. *Cortex*, 9, 82-93.
- Treisman, A. M., and Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Tyler, H. R. (1968). Abnormalities of perception with defective eye movements (Balint's syndrome). *Cortex*, 4, 154-171.
- Ullman, S. (1989). Aligning pictorial descriptions: an approach to object recognition. *Cognition*, 32, 193 - 254.

- Ungerleider, L. G. and Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (Eds.), *The Analysis of Visual Behavior*. Cambridge, MA: MIT Press.
- Van Essen, D. (1985). Functional organization of primate visual cortex. In A. Peters and E. G. Jones (Eds.), *Cerebral Cortex (Vol. 3)*. New York: Plenum Press.
- Van Essen, D. C., and Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in visual cortex. *Trends in NeuroScience*, September, 370 - 375.
- Warrington, E. K., and James, M. (1967). Disorders of visual perception in patients with localized cerebral lesions. *Neuropsychologia*, 5, 253-266.
- Warrington, E. K., and Rabin, P. (1970). Perceptual matching in patients with cerebral lesions. *Neuropsychologia*, 8, 475-487.
- Williams, M. (1970). *Brain Damage and the Mind*. Penguin Book: Middlesex, England.
- Yarbus (1967). *Eye Movements and Vision*. New York: Plenum Press.
- Yin, T.C.T., and Mountcastle, V. B. (1977). Visual input to the visuomotor mechanisms of the monkey's parietal lobe. *Science*, 197, 1381-1383.
- Zipser, D. and Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679-684.